

IN THE CLAIMS:

Please amend the claims as follows. The following listing of claims will replace all prior versions, and listings, of claims in the application.

1-31. (Canceled)

32. (New) A computer-readable memory medium storing program instructions executable to implement a method comprising:

storing a plurality of replicas of an object on each respective node of a plurality of nodes;

a first node of the plurality of nodes initiating communication with each of the remaining nodes of the plurality of nodes to attempt to synchronously update the plurality of replicas of the object, wherein the communication is successful for each node of a first subset of the remaining nodes and unsuccessful for each node of a second subset of the remaining nodes;

for each respective node of the first subset of the remaining nodes:

the respective node updating the replica of the object stored on the respective node; and

the respective node adding an identification of the object to a respective list of incoherent objects stored on the respective node in response to the communication being unsuccessful for the second subset of the remaining nodes.

33. (New) The computer-readable memory medium of claim 32, wherein the method implemented by the program instructions further comprises:

the first node updating the replica of the object stored on the first node; and

the first node adding an identification of the object to a respective list of incoherent objects stored on the first node in response to the communication being unsuccessful for the second subset of the remaining nodes.

34. (New) The computer-readable memory medium of claim 32,

wherein initiating the communication with each of the remaining nodes of the plurality of nodes comprises initiating a distributed transaction to attempt to synchronously update the plurality of replicas of the object;

wherein each node of the first subset of the remaining nodes successfully commits the transaction and wherein each node of the second subset of the remaining nodes does not successfully commit the transaction.

35. (New) The computer-readable memory medium of claim 34,
wherein at least a quorum of the plurality of nodes successfully commit the transaction.

36. (New) The computer-readable memory medium of claim 32,
wherein the first subset of the remaining nodes includes a second node of the plurality of nodes;

wherein the method implemented by the program instructions further comprises:

after the second node adding the object to the respective list of incoherent objects stored on the second node, the second node attempting to communicate with each of the other nodes of the plurality of nodes;

in response to successfully communicating with each of the other nodes of the plurality of nodes, the second node initiating a distributed transaction to synchronize the replicas stored on the second subset of the remaining nodes with the replicas stored on the first subset of the remaining nodes.

37. (New) The computer-readable memory medium of claim 36,

wherein the method implemented by the program instructions further comprises each respective node of the first subset of the remaining nodes removing the identification of the object from the respective list of incoherent objects stored on the respective node in response to said synchronizing the replicas stored on the second subset of the remaining nodes with the replicas stored on the first subset of the remaining nodes.

38. (New) The computer-readable memory medium of claim 36,

wherein each respective node of the first subset of the remaining nodes updates the replica of the object stored on the respective node by applying a first change to the replica of the object stored on the respective node;

wherein synchronizing the replicas stored on the second subset of the remaining nodes with the replicas stored on the first subset of the remaining nodes comprises each respective node of the second subset of the remaining nodes applying the first change to the replica of the object stored on the respective node.

39. (New) The computer-readable memory medium of claim 32,

wherein the object is a first object of a plurality of objects, wherein each object of the plurality of objects has a plurality of replicas stored on the plurality of nodes;

wherein the first subset of the remaining nodes includes a second node;

wherein the respective list of incoherent objects stored on the second node includes identifications of two or more of the plurality of objects;

wherein the method implemented by the program instructions further comprises:

for each respective object identified in the list of incoherent objects stored on the second node, the second node attempting to communicate with the plurality of replicas of the respective object to initiate a distributed transaction to synchronize the plurality of replicas of the respective object.

40. (New) The computer-readable memory medium of claim 32,

wherein the first subset of the remaining nodes includes a second node;

wherein the method implemented by the program instructions further comprises:

the second node periodically attempting to communicate with each of the other nodes of the plurality of nodes;

in response to determining that a particular amount of time has passed without successfully communicating with each of the other nodes of the plurality of nodes, the second node initiating an operation to create one or more new replicas of the object.

41. (New) The computer-readable memory medium of claim 32,

wherein the first subset of the remaining nodes includes a second node;

wherein the list of incoherent objects stored on the second node is stored in persistent storage of the second node;

wherein the second node adding the identification of the object to the respective list of incoherent objects stored on the second node comprises:

the second node storing information indicating addition of the identification of the object without modifying the list of incoherent objects stored in the persistent storage of the second node;

after storing the information indicating the addition of the identification of the object, the second node updating the list of incoherent objects stored in the persistent storage of the second node to reflect the addition of the identification of the object.

42. (New) The computer-readable memory medium of claim 32,

wherein each respective node updating the replica of the object stored on the respective node comprises each respective node applying a change to the replica of the object stored on the respective node.

43. (New) The computer-readable memory medium of claim 32,

wherein the plurality of replicas of the object comprises a plurality of persistent replicas of the object.

44. (New) The computer-readable memory medium of claim 32,

wherein the object is a file;

wherein each replica of the object is a replica of the file.

45. (New) The computer-readable memory medium of claim 32,

wherein the method implemented by the program instructions further comprises the first node receiving an update message;

wherein the first node initiates the communication with each of the remaining nodes in response to the update message to attempt to synchronously update the plurality of replicas of the object.

46. (New) A system comprising:

a plurality of nodes;

wherein the plurality of nodes includes memory storing program instructions executable to implement a method comprising:

storing a plurality of replicas of an object on each respective node of the plurality of nodes;

a first node of the plurality of nodes initiating communication with each of the remaining nodes of the plurality of nodes to attempt to synchronously update the plurality of replicas of the object, wherein the communication is successful for each node of a first subset of the remaining nodes and unsuccessful for each node of a second subset of the remaining nodes;

for each respective node of the first subset of the remaining nodes:

the respective node updating the replica of the object stored on the respective node; and

the respective node adding an identification of the object to a respective list of incoherent objects stored on the respective node in response to the communication being unsuccessful for the second subset of the remaining nodes.

47. (New) The system of claim 46, wherein the method implemented by the program instructions further comprises:

the first node updating the replica of the object stored on the first node; and

the first node adding an identification of the object to a respective list of incoherent objects stored on the first node in response to the communication being unsuccessful for the second subset of the remaining nodes.

48. (New) The system of claim 46,

wherein initiating the communication with each of the remaining nodes of the plurality of nodes comprises initiating a distributed transaction to attempt to synchronously update the plurality of replicas of the object;

wherein each node of the first subset of the remaining nodes successfully commits the transaction and wherein each node of the second subset of the remaining nodes does not successfully commit the transaction.

49. (New) The system of claim 46,

wherein the first subset of the remaining nodes includes a second node;

wherein the method implemented by the program instructions further comprises:

after the second node adding the object to the respective list of incoherent objects stored on the second node, the second node attempting to communicate with each of the other nodes of the plurality of nodes;

in response to successfully communicating with each of the other nodes of the plurality of nodes, the second node initiating a distributed transaction to synchronize the replicas stored on the second subset of the remaining nodes with the replicas stored on the first subset of the remaining nodes.

50. (New) A method comprising:

storing a plurality of replicas of an object on each respective node of a plurality of nodes;

a first node of the plurality of nodes initiating communication with each of the remaining nodes of the plurality of nodes to attempt to synchronously update the plurality of replicas of the object, wherein the communication is successful for each node of a first subset of the remaining nodes and unsuccessful for each node of a second subset of the remaining nodes;

for each respective node of the first subset of the remaining nodes:

the respective node updating the replica of the object stored on the respective node; and

the respective node adding an identification of the object to a respective list of incoherent objects stored on the respective node in response to the communication being unsuccessful for the second subset of the remaining nodes.

51. (New) The method of claim 50, further comprising:

the first node updating the replica of the object stored on the first node; and

the first node adding an identification of the object to a respective list of incoherent objects stored on the first node in response to the communication being unsuccessful for the second subset of the remaining nodes.

52. (New) The method of claim 50,

wherein initiating the communication with each of the remaining nodes of the plurality of nodes comprises initiating a distributed transaction to attempt to synchronously update the plurality of replicas of the object;

wherein each node of the first subset of the remaining nodes successfully commits the transaction and wherein each node of the second subset of the remaining nodes does not successfully commit the transaction.

53. (New) The method of claim 48,

wherein the first subset of the remaining nodes includes a second node of the plurality of nodes;

wherein the method further comprises:

after the second node adding the object to the respective list of incoherent objects stored on the second node, the second node attempting to communicate with each of the other nodes of the plurality of nodes;

in response to successfully communicating with each of the other nodes of the plurality of nodes, the second node initiating a distributed transaction to synchronize the replicas stored on the second subset of the remaining nodes with the replicas stored on the first subset of the remaining nodes.